

Probabilistic Model Checking for Complex Cognitive Tasks

– A case study in human-robot interaction –

Sebastian Junges¹, Nils Jansen², Joost-Pieter Katoen¹, Ufuk Topcu²

¹ RWTH Aachen University, Germany

² The University of Texas at Austin, USA

Abstract. This paper proposes to use probabilistic model checking to synthesize optimal robot policies in multi-tasking autonomous systems that are subject to human-robot interaction. Given the convincing empirical evidence that human behavior can be related to reinforcement models, we take as input a well-studied Q-table model of the human behavior for flexible scenarios. We first describe an automated procedure to distill a Markov decision process (MDP) for the human in an arbitrary but fixed scenario. The distinctive issue is that – in contrast to existing models – under-specification of the human behavior is included. Probabilistic model checking is used to predict the human’s behavior. Finally, the MDP model is extended with a robot model. Optimal robot policies are synthesized by analyzing the resulting two-player stochastic game. Experimental results with a prototypical implementation using PRISM show promising results.

1 Introduction

Verification and design for autonomous systems that work with humans account for the human’s capabilities, preferences and limitations by embedding behavioural models of *humans*. With increasing capabilities to monitor humans in dynamic, possibly mixed-reality environments, data-driven modeling enables to encode such data into behavioural models. Moreover, *reinforcement learning* [1] (RL) sufficiently describes quantitative aspects of human behaviour when solving complicated tasks in realistic environments [2,3]. Basically, RL comprises algorithms addressing the optimal control problem through learning, i.e., an agent—the human—learns how to solve a task based on repeated interaction with an environment. So-called *Q-tables* store quantitative information about possible choices of the human. We consider this information the *data* describing human behaviour.

Consider the *visio-motor* setting from [4] in Fig. 1. A human walks down a sidewalk and shall attend to three modular tasks: avoiding obstacles (purple), approaching targets (blue), or following the walkway (grey). RL




Fig. 1. Visio-motor setting with different tasks while walking down a sidewalk shown in a simulation environment. The state space is divided into the modular tasks (i) approach targets (purple), (ii) avoid obstacles (blue), and (iii) follow walkway (grey). This picture is taken from [4] with permission from Ballard.

generates Q-tables for the individual tasks; each table quantifies available choices that ultimately lead to completing the task. To build an accurate and general model of observed human behaviour for different tasks, *inverse reinforcement learning* (IRL) [5] assigns *weights* describing preferences over these tasks. It might, e. g., be more important to avoid an obstacle than to approach a target. In particular, IRL connects data sets about human behaviour—either observed over time or obtained by RL—to a general model describing how a human typically behaves in presence of different tasks. A large class of human behavioural models is covered by a set of Q-tables together with weights obtained by the methods in [4]. The given tables describe behaviour for generic scenarios in the sense that they take into account distances to features such as obstacles or targets instead of their concrete position.

This paper proposes *probabilistic model checking* [6] to analyse human behaviour models described by weighted Q-tables³ For an arbitrary *concrete scenario*, a Markov decision process (MDP) [8] is generated automatically, see Fig. 2. Intuitively, this MDP describes all possible human behaviours for this concrete scenario based on the behavioural model. The distinctive issue is that—in contrast to existing models—under-specification of the human behaviour is included. We assess the performance of the human for the scenario as well as properties of the human model itself by employing MDP model checking as supported by PRISM [9], StORM, and *iscasMc* [10]. We then assess robot behaviour in the presence of possibly disturbing humans. This is done by combining the human model with a robot-MDP model. The joint human-robot interaction model is a stochastic two-player game (SG) [11]. We synthesise optimal policies for the robot under the human behaviour using SG model-checking with PRISM-Games [12].

³ A high-level conceptual view in the form of an extended abstract is given in [7]. All technical content, implementation details, experiments and lessons learned are novel.

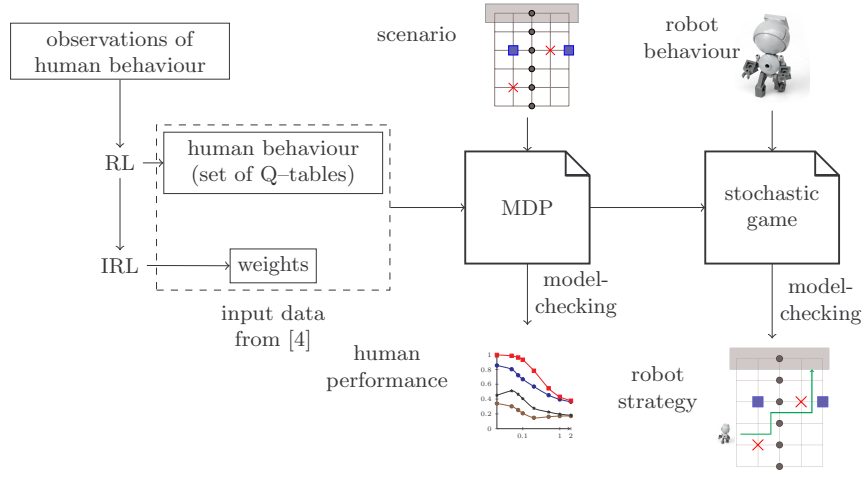


Fig. 2. Overview of our approach to verify human-robot interaction.

We stress that our approach is applicable to any human behaviour models described by weighted Q-tables. The approach is evaluated on an existing model by Rothkopf *et al.* for visio-motor tasks [13,4,14]. The main bottleneck is to handle the textual MDP description of the human behaviour of over 80,000 lines of PRISM-code. We thoroughly analysed MDPs of 10^6 states induced by a 20×20 grid scenario. The human-interaction model for a 8×8 grid scenario constitute a SG with $1.6 \cdot 10^7$ states, its generation—in absence of a symbolic engine [15,16] for SGs—takes over twenty hours; analysing maxmin reachability probabilities takes three hours. The SGs have a noticeably more complex structure than benchmarks in [17] and offer new challenges to probabilistic verification.

2 Preliminaries

In this section, we give a short introduction to models, specifications, and our notations; for details we refer to [18, Ch. 10].

Definition 1 (Probabilistic models). A stochastic game (SG) is a tuple $\mathfrak{M} = (S, s_I, Act, \mathcal{P})$ with a finite set S of states such that $S = S_\circ \uplus S_\square$, an initial state $s_I \in S$, a finite set Act of actions, and a transition function $\mathcal{P}: S \times Act \times S \rightarrow [0, 1]$ and $\sum_{s' \in S} \mathcal{P}(s, \alpha, s') \in \{0, 1\} \quad \forall s \in S, \alpha \in Act$.

- \mathfrak{M} is a Markov decision process (MDP) if $S_\circ = \emptyset$ or $S_\square = \emptyset$.
- MDP \mathfrak{M} is a Markov chain (MC) if $|Act(s)| = 1$ for all $s \in S$.

We refer to MDPs by \mathcal{M} . SGs are two-player stochastic games with players \circ and \square having states in S_\circ and S_\square , respectively. Players *nondeterministically* choose an action at each state; successors are determined *probabilistically* according to transition probabilities. MDPs and MCs are one- and zero-player stochastic games, respectively. As MCs have one action at each state, we omit this action and write $\mathcal{P}(s, s')$. For analysis, w.l.o.g. we assume that in each state there is at least one action available.

Probabilistic models are extended with *rewards* (or costs) by adding a *reward function* $\text{rew}: S \rightarrow \mathbb{R}_+$ which assigns rewards to states of the model. Intuitively, the reward $\text{rew}(s)$ is earned upon leaving the state s . Nondeterministic choices of actions in SGs and MDPs are resolved *schedulers*; here it suffices to consider memoryless deterministic schedulers [19]. Resolving all nondeterminism in SGs or MDPs yields *induced Markov chains*; note that for SG we need individual schedulers for each player.

As specifications we consider *reachability* and *expected reward* properties. A reachability property asserts that a set $T \subseteq S$ of target states is to be reached from the initial state with probability at most $\lambda \in [0, 1]$, denoted $\mathbb{P}_{\leq \lambda}(\Diamond T)$. Analogously, an expected reward property bounds the expected reward of reaching T by a threshold $\kappa \in \mathbb{R}$, denoted $\mathbb{E}_{\leq \kappa}(\Diamond T)$.

For MDPs and SGs, properties have to hold for all possible schedulers on the corresponding induced MCs. Verification can be performed by computing maximal (or minimal) probabilities or expected rewards to reach target states using standard techniques, such as linear programming, value iteration, or policy iteration [8].

3 Description of the cognitive model

This section describes the scenario of the cognitive model used as case study. It also define a formalisation that can be applied to similar case studies. This provides the basis to obtain the underlying representation as an MDP and—incorporating robot movement—an SG.

General scenario. We consider a scenario involving a human agent going over a sidewalk encountering *objects* like *obstacles* and *litter*, see Fig. 1. The human is given three modular objectives: while FOLLOW a sidewalk (represented as line) to get to the other side, she should AVOID walking into obstacles and aim to COLLECT litter. The line to follow is abstracted as a set of *waypoints*. Waypoints and objects such as obstacles and litter are called *features* of the scenario. As obstacles can be run over, litter can be collected and waypoints can be visited, these features are either being *present* or *disappeared*. For analysis purposes, we mark specific regions of

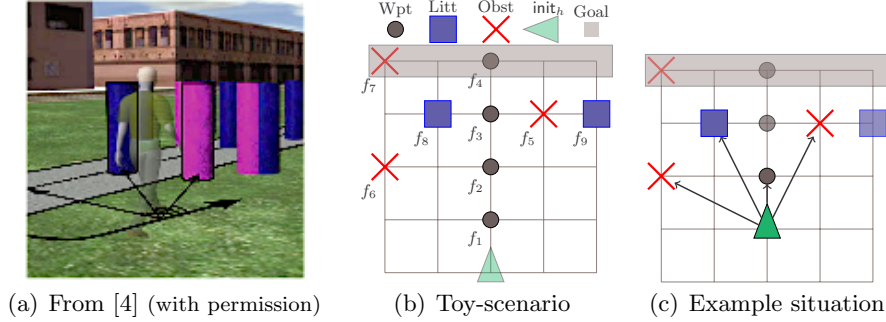


Fig. 3. Graphical representation of the grid worlds.

the environment as *goal-areas*. We assume that all features are initially present and that disappeared features remain absent. A toy-scenario is given in Fig. 3(b). The general scenario is abstracted to a discrete state model to allow to characterise human behaviour in the model. It consists of two parts: a (mostly) static environment and the dynamic human. The behaviour is determined by the movement-values assigned to each movement. These values are then translated into a probability distribution over the movements⁴.

3.1 Formal model

Let us describe how the human interacts with its environment. The environment consists of a two-dimensional grid and its features, where features have a type in $Tp = \{\text{Obst}, \text{Litt}, \text{Wpt}\}$.

Definition 2 (Environment). An environment $Env = (Loc, Feat)$ consists of a finite set of locations Loc with

$$Loc = \{(x, y) \mid x \in [0, Grid_x] \ y \in [0, Grid_y]\} \quad \text{for } Grid_x, Grid_y \in \mathbb{N},$$

and a set of features $Feat \subseteq Tp \times Loc$. A feature $f = (tp_f, \ell_f) \in Feat$ consists of a type and a (feature-)location.

Features are partitioned according to their type: $Feat = Feat_{\text{Obst}} \cup Feat_{\text{Litt}} \cup Feat_{\text{Wpt}}$, such that $Feat_{tp} \subseteq \{tp\} \times Loc$.

Example 1. Consider our running example in Fig. 3(b). The environment is given as $Env = (Loc, Feat)$ with $Loc = \{(x, y) \mid x \in [0, 4] \ y \in [0, 4]\}$, and

$$Feat = \{f_i = (\text{Wpt}, (2, i)) \mid i \in \{1 \dots 4\}\}$$

⁴ We use the term movement to avoid confusion with actions in MDPs.




Fig. 4. Orientations and human movement.

$$\begin{aligned} &\cup \{f_5 = (\text{Obst}, (3, 3)), f_6 = (\text{Obst}, (0, 2)), f_7 = (\text{Obst}, (0, 4))\} \\ &\cup \{f_8 = (\text{Litt}, (1, 3)), f_9 = (\text{Litt}, (4, 3))\} \end{aligned}$$

Human. The human h is represented by its *position* $\text{pos}_h = (\ell_h, \alpha_h)$ which is a pair of a location ℓ_h and an orientation α_h . An orientation has 8 possible directions, i.e. $\alpha_h \in \text{Orient} = \{i \cdot \frac{1}{4}\pi \mid i \in [0, 7]\}$. We assume that the human starts in position init_h .

Human movement. For each direction, let the associated direction vector $\text{Dir}: \text{Orient} \rightarrow \{-1, 0, 1\}^2 \setminus \{(0, 0)\}$, see Fig. 4(a). Human movements $M_h = \{\text{LEFT}, \text{STRAIGHT}, \text{RIGHT}\}$, see Fig. 4(b), have associated changes in angle of $\beta = -\frac{1}{4}\pi, 0$, or $\frac{1}{4}\pi$. To be precise, the *post-position* $\text{post}_m(\text{pos}_h)$ for human position $\text{pos}_h = (\ell_h, \alpha_h)$ and movement m with β_m , is given by:

$$\text{post}_m(\text{pos}_h) = (\ell_h + \text{Dir}(\alpha_h + \beta_m), (\alpha_h + \beta_m) \bmod 2\pi).$$

A movement m is *valid in* pos_h if $\text{post}_m(\text{pos}_h)$ is a well-defined position, that is, the location of the post-position is on the grid.

Situations. The status of a scenario is called a *situation*:

Definition 3 (Situation). Let $\text{Env} = (\text{Loc}, \text{Feat})$ be an environment. A situation s is a pair of a human location and the present features $(\text{pos}_h, \text{PFeat}) \in \text{Loc} \times 2^{\text{Feat}}$.

As we see in Sect. 3.2, the situations define the state space of our model. If the post-position of a movement coincides with a feature, then the feature is marked disappeared. Formally, for a valid movement m in position pos_h , the present features PFeat are updated as follows:

$$\text{fu}_m(s) = \text{PFeat} \setminus \{(\text{tp}, \text{post}_m(\text{pos}_h))\} \subseteq \text{Feat}.$$

The effect of a movement of the human is as follows.

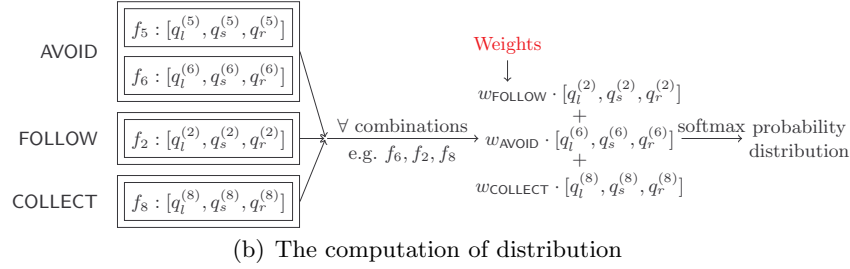



Fig. 5. Sketched procedure to construct sets of distributions.

Definition 4 (Human movement effect). *The effect of a human movement m in a situation $s = (\text{pos}_h, \text{PFeat})$ is the situation*

$$\text{eff}_m(s) = (\text{post}_m(\text{pos}_h), \text{fu}_m(s)).$$

Distance and angle. Let f be a feature and $\text{pos}_h = (\ell_h, \alpha_h)$ the position of the human h . The distance between h and f is given by the Euclidean norm, i.e. $d_h(f) = \|\ell_h - \ell\|_2$, and $\gamma_h(f)$ denotes the signed angle between the human orientation and $\ell_h - \ell$.

Example 2. In the situation $s = (\text{pos}_h, \text{PFeat})$ depicted in Fig. 3(c), $\text{eff}_{\text{LEFT}}(s) = (((1, 2), \frac{3}{4}\pi), \text{PFeat} \setminus \{f_2\})$. The distance of the human h to f_9 at $(4, 3)$ is $d_h(f_9) = \|(4, 3) - (2, 1)\|_2 = 2\sqrt{2}$, the angle is $\gamma_h(f_9) = -\frac{1}{4}\pi$.

Movement-valuation. Let us consider how to obtain movement values when combining objectives. These values are a kind of importance factor of the movement in a particular situation. The steps are outlined in Fig. 5. Fig. 5(a) outlines how to obtain values for the possible movements for each objective. These values are not always unique. Fig. 5(b) describes how to combine the movement values for different objectives, and to translate it into distributions. Details of the various steps are given below.

Relevant features. Recall that we consider three objectives, i.e. $\text{Objectives} = \{\text{FOLLOW}, \text{AVOID}, \text{COLLECT}\}$. For each objective o , we have exactly one

corresponding feature-type $f(o)$; Waypoints, obstacles and litter correspond to FOLLOW, AVOID and COLLECT, respectively. As the behaviour of the human is independent of disappeared features, we call the set of present features the *relevant features* $\text{RelFeat}_o(s) = \{(f(o), \ell) \in \text{PFeat}\}$ w.r.t. objective o in situation $s = ((\ell_h, \alpha_h), \text{PFeat})$.

Closest relevant features. We adopt the assumption from [4] that for each objective o , only the closest feature of type $f(o)$ is relevant for the behaviour with respect to o .

Remark 1. While this assumption is strong, it reduces the number of hidden parameters a learning method has to estimate. Our method can be easily adapted to models where other/multiple features per type are relevant.

Let $\text{Close}(s, o)$ be *closest relevant features* for objective o in situation s :

$$\text{Close}(s, o) = \{f \in \text{RelFeat}_o(s) \mid \forall f' \in \text{RelFeat}_o(s). d_h(f) \leq d_h(f')\}.$$

Example 3. For the situation s depicted in Fig. 3(b), $\text{Close}(s, \text{AVOID}) = \{f_5, f_6\}$ and $\text{Close}(s, \text{FOLLOW}) = \{f_1\}$.

While in the physical reality two objects are almost never equally far away, in the grid abstraction, this happens frequently. Thus, the set of closest relevant features is not necessarily a singleton. The human behaviour is underspecified in this case, and any of the features in the set might be the actually relevant feature.

While we support under-specification, it can be valuable to collapse this, and create the *unique* closest relevant features. We do so by selecting the feature from the closest relevant features with (i) the smallest absolute angle, and in case of a tie (ii) left of the human over right-of-the human. This assumption allows us to treat larger benchmarks at the cost of a less precise model.

Example 4. Breaking the tie means that in the situation s depicted in Fig. 3(b), $\text{Close}(s, \text{AVOID}) = \{f_5\}$, as $|\gamma_h(f_5)| \approx 26^\circ < |\gamma_h(f_6)| \approx 63^\circ$.

Movement-values. We assume that we have Q-tables as in [4]. In particular, for each *movement* m and each objective o , a Q-table $Q_o^m: \mathbb{R} \times \mathbb{R}_+ \rightarrow \mathbb{R}$ maps the angle $\gamma_h(f)$ and distance $d_h(f)$ between a human and a close relevant feature f to an *objective-movement-value*. Partial tables are given in Tab. 1. Alike to [4], we translate the set of closest relevant features

Table 1. Illustrative fragment of Q-tables $Q_o^m(\gamma, d)$ with the bins for angle γ_h (in degrees) and distance d_h (in tiles) listed vertically and horizontally, respectively. Boldface entries correspond to the lowest (avoiding the movement the most) value over the different actions, while the coloured entries are used in Ex. 5.

(a) AVOID, LEFT					(b) AVOID, STRAIGHT					(c) AVOID, RIGHT				
	[0, 1]	(1, 2]	(2, 3]	...		[0, 1]	(1, 2]	(2, 3]	...		[0, 1]	(1, 2]	(2, 3]	...
...
$[-90, -45]$	-0.5	-0.495	-0.01	...	$[-90, -45]$	0	0	0	...	$[-90, -45]$	-0.96	-0	0	...
$[-45, -15]$	-1.21	-1.19	-0.59	...	$[-45, -15]$	-0.02	-0.96	-0.28	...	$[-45, -15]$	-0.88	-0.77	0	...
$[-15, +15]$	-0.87	-1.12	-0.43	...	$[-15, +15]$	-0.4	-1	-1.12	...	$[-15, +15]$	-0.88	-1.3	-0.04	...
$[+15, +45]$	-0.75	-0.18	0	...	$[+15, +45]$	-0.27	-0.99	-0.67	...	$[+15, +45]$	-0.99	-0.99	-0.32	...
$[+45, +90]$	-0.75	-0.06	0	...	$[+45, +90]$	0	0	0	...	$[+45, +90]$	-0.5	-0.13	-0.34	...
...

into a set of objective-movement-vectors $V^o(s)$ by a lookup in the Q-table, and store the feature for later use.

$$V^o(s) = \{(f, [q_l, q_s, q_r]) \mid f \in \text{Close}(s, o), \quad q_l = Q_o^{\text{LEFT}}(\gamma_h(f), d_h(f)), \\ q_s = Q_o^{\text{STRAIGHT}}(\gamma_h(f), d_h(f)), \quad q_r = Q_o^{\text{RIGHT}}(\gamma_h(f), d_h(f)) \}$$

The vector entries collect movement-values with respect to a fixed feature.

Example 5. For s as depicted in Fig. 3(b), $\text{Close}(s, \text{AVOID}) = \{f_5, f_6\}$, we obtain by lookup in Tab. 1 – using $d_h(f_i) = \sqrt{5} \in [2, 3]$ for $i \in \{5, 6\}$, $\gamma_h(f_5) \approx 26^\circ \in [15, 45]$, and $\gamma_h(f_6) \approx -63^\circ \in [-90, -45]$ – the following result: $V^{\text{AVOID}}(s) = \{(f_5, [0, -0.67, -0.32]), (f_6, [-0.01, 0, 0])\}$.

Combining movement values. We assume that we have an *objective-weight vector* $w = [w_{\text{AVOID}}, w_{\text{COLLECT}}, w_{\text{FOLLOW}}] \in \text{Distr}(\text{Objectives})$ over the objectives, e.g. obtained by IRL as in [4]. The objective-movement-vectors are translated into movement-vectors by calculating a weighted sum over all combinations: We first scale all $V^o(s)$ with w_o , that is, the weighted-objective-movement values $wV^o(s) = \{(f, w_o \cdot \mathbf{x}) \mid (f, \mathbf{x}) \in V^o(s)\}$. Then, we take the sum of the weighted vectors and construct the union of all involved features. Formally, for any movement-vector and position of the human the corresponding set of movement-values, a pair consisting of the *movement-vector* and a set of the features involved:

$$V(s) = \{(\{f, f', f''\}, \mathbf{x} + \mathbf{x}' + \mathbf{x}'') \mid (f, \mathbf{x}, f', \mathbf{x}', f'', \mathbf{x}'') \in \bigtimes_{o \in \text{Objectives}} wV^o(s)\}$$

where the sum is the component-wise sum.

Example 6. Similar to Ex. 5, we obtain $V^{\text{COLLECT}}(s) = \{(f_8, [2.5, 2.4, 2])\}$, $V^{\text{FOLLOW}}(s) = \{(f_1, [0, 1.14, 0])\}$. The $V(s)$ can be computed using the objective-weights (provided by e.g. IRL) $\mathbf{w} = [0.414, 0.215, 0.369]$. We get

$$\begin{aligned} V(s) &= \{(\{f_6, f_2, f_9\}, 0.414 \cdot [\dots, \dots, \dots] + 0.215 \cdot [\dots, \dots, \dots] + 0.369 \cdot [\dots, \dots, \dots]), \\ &\quad (\{f_7, f_2, f_9\}, 0.414 \cdot [\dots, \dots, \dots] + 0.215 \cdot [\dots, \dots, \dots] + 0.369 \cdot [\dots, \dots, \dots])\} \\ &= \{(\{f_6, f_2, f_9\}, [0.53, 0.66, 0.3]), (\{f_7, f_2, f_9\}, [0.53, 0.93, 0.42])\} \end{aligned}$$

Notice that at the grid borders, only some movements are possible. Thus, we need to rule out such movements. In the course of this paper, we do this by resetting those movement-values to $-\infty$. If no movement is possible, we remove the vector from the movement values.

Distribution. The last step (cf. Fig. 5(b)) is to transfer movement values into a distribution. Under the working hypothesis that the valuation reflects the likelihood that a human makes a specific move, we can translate this into a distribution over the movements – provided any movement is possible.

The stochastic behaviour of the human is obtained by translating any vector $\mathbf{x} \neq -\infty$ for $(F, \mathbf{x}) \in V(s)$ for some situation s to a distribution over movements, by means of a *softmax*-function $[1]: \mathbb{R}_{\infty}^n \rightarrow [0, 1]^n$ – which attributes most, but not all probability to the maximum, hence the name. Using $e^{-\infty} := 0$, the distribution is defined as:

$$\text{softmax}_{\tau}(\mathbf{x})_i = \frac{e^{x_i/\tau}}{\sum_{i \leq |\mathbf{x}|} e^{x_i/\tau}}.$$

For any invalid movement, the denominator is thus unaffected as the term is zero, the numerator is zero, effectively ruling out the transition. The parameter τ is called the *temperature*. Towards a zero temperature, the function is as a (hard) maximum, while with a high temperature it yields an almost uniform distribution.

3.2 MDP model of the human behaviour

Given the formal description above, we are now ready to construct the MDP for the human behaviour. The state space is given by the set of situations. The initial state is given by the start-location of the human and the assumption that initially all features are present. As stated before we have two possible sources of non-determinism: (1) Under-specification of the model: the closest relevant feature is not unique. (2) Insufficient confidence in some entries in the Q-table, e.g. if the amount of data does

not allow us to draw conclusions. Here we only consider the first source, the latter is an extension with some more non-determinism.

Non-determinism due to under-specification. Resolving the non-singleton movement-vectors is modelled to be an action of the environment.

Definition 5. *The MDP $\mathcal{M} = (S, s_I, Act, \mathcal{P})$ reflecting the human behaviour starting in $init_h$ on a environment $Env = (Loc, Feat)$ using temperature τ is given by*

- $S = \{(pos_h, P) \mid pos_h = (\ell_h, \alpha_h) \in Loc \times Orient, P \subseteq Feat\}$
- $s_I = (init_h, Feat)$
- $Act = Feat^3$
- $\mathcal{P}(s, a) = \begin{cases} \{eff_{(M_h)_i}(s) \mapsto \text{softmax}_\tau(\mathbf{x})_i \mid i \in \{1, 2, 3\}\} & (a, \mathbf{x}) \in V(s) \\ 0 & \text{otherwise.} \end{cases}$

Rewards for the human performance. In order to evaluate the performance with respect to the objectives, we define *transition-reward* mappings rew_o for each $o \in \text{Objectives}$,

$$\forall s, s' \in S \forall a \in Act, \text{rew}_o(s, a, s') = \begin{cases} 1 & \text{if RelFeat}_o(s) \neq \text{RelFeat}_o(s') \\ 0 & \text{otherwise.} \end{cases}$$

To define a combined reward, we want to avoid obstacles and collect litter, which means to penalise visiting some fields while rewarding others; this calls for the use of both positive and negative rewards, which is not possible in PRISM. With the help of the goal-states, we can give rewards upon entering them and considering the performance afterwards⁵.

3.3 SG model of human-robot interaction

In order to obtain a model which models human behaviour and allows us to synthesise a plan for the robot, we have to consider a unified model. As we want to choose actions for the robot, the robot is naturally modelled as a (potentially non-probabilistic) MDP. Notice that the non-determinism of the robot is *controllable*, whereas the non-determinism of the human model is *uncontrollable*. This naturally leads to a stochastic two-player game. As a design choice, we let robot and human move in turns (typically alternating). While this abstraction is not inherently different from synchronous movements, this means that we can use single actions to determine the movements.

⁵ In combination with the robot, this prevents distinguishing robot and human performance: here encoding the reward in the state space would be our last resort.

The robot model. We can support any MDP on the same grid given as a PRISM-module. Within the scope of the paper, we considered a simple robot either turning 90-degrees (left or right) in place, or moving forward. Synchronisation with the environment is via shared variables.

Considering the robot movement. We both (1) considered a scenario where the human behaviour is not influenced by the robot and (2) the robot as an obstacle-feature. (1) yields a controller for the robot which is not intrusive, whereas (2) means that the human tries to avoid the robot and the controller takes this into account. For the latter, we had no suitable data and we assumed that the human treats the robot as if the robot were static. However, the methodology presented here can be used as is if suitable Q-tables for dynamic obstacles are available.

4 Experiments

We developed a prototypical tool-chain to show the feasibility of our approach outlined in Fig. 2. The tool realizes Fig. 5 for any scenario and exploits the PRISM-Games v2.0beta3 [12]. We first discuss how to create the model and then present some evaluation results obtained under Debian 8 on a HP BL685C G7, 48 cores, 2.0GHz each, and 192GB of RAM.

State and transition encoding. We encode the SG for the joint human and (optional) robot behaviour in the PRISM-language. The model has a module for the human and one for the robot. A global flag indicates whether the human (or robot) moves next. The robot (human) has a precondition that the robot (human) may move and an update to let the human (robot) move next. A module for the human consists of 3 integers to represent her location and orientation, and a boolean b_f for each feature f , with b_f true iff f is present. The location of the (static) features are constants. Though the set of reachable scenarios is exponential in the number of features, the encoding is cubic as the behaviour is based on the nearest present features only; see also App. A.

Reward encoding. The PRISM language does not support state-action-target rewards as used in Sect. 3.3. We support two options for encoding the objective-reward: (1) Rescale the rewards to state-action rewards (preserving for expected reward measures [8]). As the rescaling depends on the probabilities, we de-facto have to scale rewards for each command which may reach a feature separately. Commands can only be named in the scenario without a robot; the combination of named commands and global variables (the status of the features) is not supported. (2) Introduce a first-time flag for any location with features indicating whether it is visited

for the first time. This increases the number of states by $|\{\ell \mid (\text{tp}, \ell) \in \text{Feat}\}| \cdot |\text{Orient}|$. The number of commands does not increase, the BDD size is hardly affected. The rewards are attached to the states: the encoding is as large as the number of locations with features.

Optimisations. We investigated various performance improvements. Two notable effective insights are: (1) (Only) the Q-table for obstacle avoidance shows equal values for the far-away bins—indicating that human behaviour does not consider far-away obstacles. It is thus not necessary to distinguish which obstacle is nearest once they all induce the same lookup-value. This is especially relevant on large and sparse scenarios. (2) As every human location occurs on a large number of commands, naming expressions (called formulas in the PRISM-language) reduces the overhead of specifying the location repeatedly. Using short non-descriptive variable names reduces the encoding further. Together, they reduce the parsing time by 40%.

Model construction and parsing. As most case studies can be succinctly described in the PRISM-language, parsing usually is not an issue. Here it is. Our models are up to 100,000 lines of PRISM code. App. B lists details about the model sizes and their building times. Parsing takes a significant amount of time. The hybrid PRISM engine yields good overall performance: model construction is much faster than for explicit state spaces, and value iteration with many small probabilities (< 0.01) takes many iterations to converge, which is slow on a purely BDD-based engine.

Evaluating the human. For the MDP model w/o robot, we compute some minimum and maximum probabilities; their difference indicates the relevance of the underspecification. Fig. 6(a) gives some verification results for a 20x20 grid with 2 landmarks, 2 obstacles, and 7 waypoints. The module contains 84,000 commands. Fig. 6(b) plots the min/max probability to reach the goal area when the human is only told to follow the waypoints against the temperature (controlling the variability in the softmax function). It shows that with low variability most humans indeed reach the other side without leaving the grid. This quickly drops with higher variability (where any features are mostly ignored). Fig. 6(c) indicates a similar behaviour for step-bounded reachability for different number of steps (x-axis) and temperatures 0.05 and 0.5. Most humans need > 30 steps to reach the goal, indicating that based on the given data, they very unlikely walk in straight lines. The gap due to underspecification is significant as long as the variability is not too high. With low variability, most humans arrive within 60 steps. Detailed analyses considering the obtained schedulers show where underspecification has the largest effect.

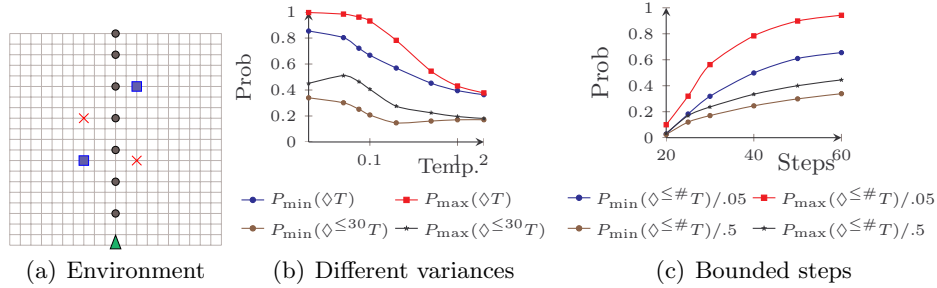


Fig. 6. Human performance evaluation.

Human-robot interaction. The state space is about a factor robot locations times robot directions (4) times turn-flag (2) larger. The number of choices is roughly twice the number of states. We built a 8×8 grid with 2/2/3 features; construction of the $1.6 \cdot 10^7$ states took 20 hours. (Using the explicit engine, as **PRISM-Games** has no symbolic engine.) Model checking of a min/max reachability probability took 3 hours. Constructing a 11×11 grid with 2/2/3 features with $7.4 \cdot 10^7$ SG-states took 198 hours. States are slowly created compared to typical benchmarks due to the large number of commands, rendering a thorough analysis impossible. We made simplifications: (1) resolve the underspecification to have one player making trivial choices, or (2) consider a coalition where the underspecification is controllable. Both observations allow us to do a (suboptimal) analysis on the (isomorphic) single-player SG, i.e. an MDP. While this does not reduce the size of the model, it allows for using symbolic engines. We give figures for the coalition approach here; removing underspecification was slightly faster. The aforementioned 11×11 grid-model takes 191,290 MTBDD nodes (13,725 terminal) in the symbolic engine. Model construction in 19 seconds. Checking unbounded reachability takes about 500 iterations and four hours. As we are interested in the robot reaching its goal, finite-horizon plans are relevant. The first 100 iterations are done within 40 minutes and yield an almost identical scheduler.

5 Conclusion and discussion

We have successfully translated a cognitive model into a formal setting and used it to compute control plans for robots moving in the presence of humans handling complex tasks. In the paper, we discussed the model as well as several (open) challenges we stumbled upon.

Lessons learned: the model. Based on the model-checking results, we obtained five lessons about the weighted Q-table model. (1) Fig. 6(c)

shows that humans most likely walk in wavy lines. This is due to the lack of a *notion of progress* in visiting waypoints—it does not penalise walking in circles, as only positive reward is earned on visiting waypoints. Following a line and giving a penalty for any diverging move (as in [4]) would provide a notion of progress. (2) As the Q-tables do not *take into account the border of the environment*, they are not avoiding a deadlock (or unspecified behaviour): The probability for leaving the grid is substantial in many cases. (3) For the discrete model, there is a potentially huge difference in behaviour based on how the underspecification is resolved; *any analysis on the learned model has to take this underspecification into account*. (4) Modelling variability over human behaviour by a single softmax and *using a memoryless model* are rough estimates. Therefore it is quite likely that a human behaves inconsistent (in the SG) to statements in e.g. [14]. (5) Finally, the Q-tables contain some unexpected *outliers* which in some configurations lead to unexpected behaviour.

Lessons learned: the method. The description for the human behaviour including variability can be translated into a formal model; allowing a variety of properties to be easily analysed—in particular, it allows for analysing under-specifications and ill-defined data. The generalisation to a robot planning scenario (and to a SG) is straightforward, and enables to compute plans fulfilling specific properties. Probabilistic verification of this model raised *six challenges*: (1) The *large number of different probabilities* in learned data blows up the encoding (and the BDD), it also prevents successful application of typical reduction techniques such as bisimulation. It would be interesting to *regularise the model* on the learning side, or use techniques like *ϵ -bisimulation* [20]. (2) The softmax function serves the only purpose of introducing variability; its nice features w.r.t. differentiability are not relevant here. *Sensitivity analysis over Q-table values would be of interest* but current parameter synthesis techniques based on rational functions [21,22] cannot cope with exponentials. (3) Despite the lacking of typical symmetries in the scenarios, its encoding is significantly smaller than enumerating all states, as (here) the behavior only depends on the nearest obstacles: even for three features of every kind, reduction of a factor over 50 results. (4) While the approach yields promising results for the MDP scenario, the SG suffers from a state space explosion. Although the turn-based game is highly regular, this cannot be exploited by *the lack of a symbolic engine* to solve SGs. (5) The explicit engine suffers from the *loss of information* in the PRISM-encoding: formulas grouping states and making explicit that in a set of states only one or two commands are relevant are lost. (6) Finally, the scheduler is always computed for the

full state space, whereas based on reachability of many states, *we are only interested in a fragment of the full state space*; for states with a low reachability probability, we can take any action.

Future work. For future work, we would also like to investigate automatic abstraction techniques as in [23] and restrict exploration of the model as in [24], as well as sensitivity analysis and/or model repair, potentially based on techniques presented in [22] or those in [25].

Acknowledgement. We thank Mary Hayhoe and Matthew Tong for providing and explaining the data.

References

1. Sutton, R.S., Barto, A.G.: Reinforcement Learning: An Introduction. MIT Press (1998)
2. Barto, A.G.: Adaptive critics and the basal ganglia. In: Models of Information Processing in the Basal Ganglia. MIT Press (1995) 215–232
3. Daw, N.D., Doya, K.: The computational neurobiology of learning and reward. *Current Opinion in Neurobiology* **16**(2) (2006) 199–204
4. Rothkopf, C.A., Ballard, D.H.: Modular inverse reinforcement learning for visuomotor behavior. *Biological Cybernetics* **107**(4) (2013) 477–490
5. Ng, A.Y., Russell, S.J.: Algorithms for inverse reinforcement learning. In: Proc. of ICML, Morgan Kaufmann (2000) 663–670
6. Katoen, J.P.: The probabilistic model checking landscape. In: Proc. of LICS. (2016) 31–46
7. Junges, S., Jansen, N., Katoen, J.P., Topcu, U.: Probabilistic verification for cognitive models. In: Proc. of Cross-Disciplinary Challenges for Autonomous Systems (CDCAS). AAAI Fall Symposium (2016)
8. Puterman, M.L.: Markov Decision Processes: Discrete Stochastic Dynamic Programming. John Wiley and Sons (1994)
9. Kwiatkowska, M., Norman, G., Parker, D.: PRISM 4.0: Verification of probabilistic real-time systems. In: Proc. of CAV. Volume 6806 of LNCS, Springer (2011) 585–591
10. Hahn, E.M., Li, Y., Schewe, S., Turrini, A., Zhang, L.: iscasMc: A web-based probabilistic model checker. In: Proc. of FM. Volume 8442 of LNCS, Springer (2014) 312–317
11. Condon, A.: The complexity of stochastic games. *Inf. Comput.* **96**(2) (1992) 203–224
12. Kwiatkowska, M., Parker, D., Wiltsche, C.: PRISM-Games 2.0: A tool for multi-objective strategy synthesis for stochastic games. In: Proc. of TACAS. Volume 9636 of LNCS, Springer (2016) 560–566
13. Rothkopf, C.A., Ballard, D.H.: Credit assignment in multiple goal embodied visuomotor behavior. *Frontiers in Psychology* **1** (2010) 217–230
14. Tong, M.H., Zohar, O., Hayhoe, M.M.: Control of gaze while walking: task structure, reward, and uncertainty. *Journal of Vision* (2016) (submitted).
15. Parker, D.: Implementation of Symbolic Model Checking for Probabilistic Systems. PhD thesis, University of Birmingham (2002)

16. Baier, C., Clarke, E.M., Hartonas-Garmhausen, V., Kwiatkowska, M.Z., Ryan, M.: Symbolic model checking for probabilistic processes. In: Proc. of ICALP. Volume 1256 of LNCS (1997) 430–440
17. Chen, T., Forejt, V., Kwiatkowska, M., Parker, D., Simaitis, A.: PRISM-games: A model checker for stochastic multi-player games. In: Proc. of TACAS. Volume 7795 of LNCS, Springer (2013) 185–191
18. Baier, C., Katoen, J.P.: Principles of Model Checking. The MIT Press (2008)
19. Vardi, M.Y.: Automatic verification of probabilistic concurrent finite-state programs. In: Proc. of FOCS, IEEE CS (1985) 327–338
20. Bacci, G., Bacci, G., Larsen, K.G., Mardare, R.: Computing behavioral distances, compositionally. In: Proc. of MFCS. Volume 8087 of LNCS, Springer (2013) 74–85
21. Hahn, E.M., Hermanns, H., Zhang, L.: Probabilistic reachability for parametric Markov models. *Software Tools for Technology Transfer* **13**(1) (2010) 3–19
22. Dehnert, C., Junges, S., Jansen, N., Corzilius, F., Volk, M., Bruintjes, H., Katoen, J.P., Abraham, E.: Prophesy: A probabilistic parameter synthesis tool. In: Proc. of CAV. Volume 9206 (2015) 214–231
23. Wachter, B., Zhang, L., Hermanns, H.: Probabilistic model checking modulo theories. In: Proc. of QEST, IEEE CS (2007) 129–140
24. Brázdil, T., Chatterjee, K., Chmelik, M., Forejt, V., Kretínský, J., Kwiatkowska, M.Z., Parker, D., Ujma, M.: Verification of Markov decision processes using learning algorithms. In: Proc. of ATVA. Volume 8837 of LNCS, Springer (2014) 98–114
25. Chen, T., Hahn, E.M., Han, T., Kwiatkowska, M., Qu, H., Zhang, L.: Model repair for Markov decision processes. In: Proc. of TASE, IEEE CS (2013) 85–92

A The size of the encoding

We argue that the encoding is cubic. For the description of the transition probabilities, it suffices to only consider under which circumstances a feature is not in $\text{Close}(s, o)$: That is exactly if there is another feature f' of type $f(o)$ s.t. $d_h(f') < d_h(f)$. For any action $[f_{\text{Obst}}, f_{\text{Litt}}, f_{\text{Wpt}}] \in \text{Feat}_{\text{Obst}} \times \text{Feat}_{\text{Litt}} \times \text{Feat}_{\text{Wpt}}$ and each human position, we have as a precondition $\bigwedge_{o \in \text{Objectives}} \bigwedge_{f \in \text{Feat}} \{\neg b_f | f \in \text{Feat}_o \wedge d_h(f_o) < f\}$. That means that for each human position we have $\leq \prod_{o \in \text{Objectives}} |\text{Feat}_{f(o)}|$ commands.

If avoiding the robot is to be treated as a separate objective, than the transition encoding for any command in the human module has to be given separately for every possible location of the robot.

B Experimental Performance

Table 2(a) gives the grid size, the number of obstacles, landmarks, and waypoints, the number of commands, the number of states, choices and branches of the underlying model using $\tau = 0.075$, the number of decision diagram nodes (with the number of terminal nodes in parentheses) as well as the time to parse and the time to build the model, either explicitly or via the symbolic engine. A – indicates that the figure could not be obtained within 20 hours. The high number of terminal mtbdd nodes is a major challenge for the dd-based engines, whereas the explicit engine cannot cope with the large number of commands.

In Table 2(b) we give the same information for the models including a robot; notice that here, the symbolic numbers where obtained by transforming the model into an MDP. Please notice that the decision diagrams are comparatively small: This is due to the regularity of the alternating movements.

For the (8,8) SG models, the model checking time to obtain a scheduler maximising the probability to reach the other side of the grid without crashing, taking into account any scheduler resolving the underspecification is roughly 1300, 3000, 5400 and 10700 seconds, respectively.

Table 2. Building times

(a) Human MDP; non-determinism due to underspecification										
grid	features	#	commands	states	choices	transitions	dd-nodes	t_{parse}	$t_{\text{build}}^{\text{exp}}$	$t_{\text{build}}^{\text{sym}}$
(11,11)	2/2/3		13325	9.4E4	2.4E5	9.9E4	131172 (13727)	7	965	11
(11,11)	4/4/4		33362	2.9E6	3.3E6	8.0E6	534607 (32540)	14	-	97
(16,16)	2/2/3		21256	1.1E5	1.2E5	3.1E5	135384 (16220)	27	2337	22
(16,16)	4/4/4		59527	3.5E6	4.0E6	1.0E7	414431 (34484)	41	-	194
(21,21)	2/2/3		42077	2.0E5	2.1E5	5.8E5	181376 (18236)	100	7192	41
(21,21)	4/4/4		95356	6.4E6	6.7E6	1.8E7	496709 (40441)	163	-	388
(21,21)	2/2/7		83900	3.2E6	3.4E6	9.3E6	466540 (30638)	135	-	310

(b) Human & Robot 2p SG; non-determinism due to underspecification										
grid	features	#	commands	states	choices	transitions	dd-nodes	t_{parse}	$t_{\text{build}}^{\text{exp}}$	$t_{\text{build}}^{\text{sym}}$
(8,8)	1/1/2		2864	2.5E6	4.8E6	6.4E6	31097 (2511)	3.0	2181	2.5
(8,8)	1/1/3		3775	4.9E6	9.5E6	1.3E7	49580(3262)	3.4	5951	3.9
(8,8)	2/1/3		4464	8.4E6	1.6E7	2.2E7	72988 (4043)	3.5	14759	5.7
(8,8)	2/2/3		6624	1.6E7	3.2E7	4.4E7	128466 (6346)	3.7	70595	10.1
(11,11)	2/2/3		13179	7.4E7	1.4E8	2.0E9	191290 (13725)	7.5	7.1E5	19.2
(16,16)	2/2/7		50335	3.1E9	6.4E9	9.5E9	911575 (22870)	41.3	-	257

These files are also available on GitHub: https://github.com/moves-rwth/human_factor_models.